Matrix Completion from a Few Entries

Raghunandan H. Keshavan*, Andrea Montanari*†, and Sewoong Oh*

August 4, 2009

Abstract

Low-rank models are frequently used in machine learning and statistics. An important domain of application is provided by collaborative filtering where one of the problems is to reconstruct a low rank matrix from a sparse subset of entries. Given M, an $n\alpha \times n$ matrix of rank $r \ll n$, assume that a uniformly random subset E of its entries are observed. We prove that a simple spectral algorithm reconstructs M from $|E| = O(r\, n)$ observed entries with arbitrarily small relative root mean square error. Further, we describe an efficient algorithm that reconstructs M *exactly* from $|E| = O(n \log n)$ entries (for bounded r) under appropriate technical conditions. We then present results to show robustness of the algorithm when the entries observed are corrupted by additive noise. In this context, we prove an upper bound on the root mean square error that is order-optimal in a number of circumstances.

*Department of Electrical Engineering, Stanford University

†Departments of Statistics, Stanford University